



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ


EP/05/1479

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Stephen Hordley

Dated

22 February 2005

Patents Act 1977
(Rule 16)

The
Patent
Office

THE PATENT OFFICE

12 FEB 2004

LONDON

1/77

The Patent Office

Cardiff Road
Newport
South Wales

13FEB04 E872960 P1082829

P0177700 0.00-0403159.7 ACCOUNT CHA

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

1. Your reference	J00046278GB		
2. Patent application number (The Patent Office will fill this part in)	0403159.7		12 FEB 2004
3. Full name, address and postcode of the or of each applicant (<u>underline all surnames</u>)	Parc Technologies Limited 8th Floor, The Tower Building 11 York Road London SE1 7NX United Kingdom		
Patents ADP number (if you know it)	8013823002		
If the applicant is a corporate body, give the country/state of its incorporation	United Kingdom		
4. Title of the invention	Traffic Flow Determination in Communications Networks		
5. Name of your agent (if you have one)	RGC Jenkins & Co.		
"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)	26 Caxton Street London SW1H 0RJ United Kingdom		
Patents ADP number (if you know it)	03966736001 ✓		
6. Priority: Complete this section if you are declaring priority from one or more earlier patent applications, filed in the last 12 months	Country	Priority application number (if you know it)	Date of filing (day / month / year)
7. Divisionals, etc: Complete this section only if this application is a divisional application or resulted from an entitlement dispute (see note f)	Number of earlier UK application		Date of filing (day / month / year)
8. Is a Patents Form 7/77 (Statement of inventorship and of right to grant of a patent) required in support of this request?	Yes		
Answer 'Yes' if: a) any applicant named in part 3 is not an inventor, or b) there is an inventor who is not named as an applicant, or c) any named applicant is a corporate body. Otherwise answer NO (See note d)			

Patents Form 1/77


9. Accompanying documents: A patent application must include a description of the invention. Not counting duplicates, please enter the number of pages of each item accompanying this form:

Continuation sheets of this form

Description 26

Claim(s) 7

Abstract 1

Drawing(s) 5 

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translation of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77) Yes (1)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11. I/We request the grant of a patent on the basis of this application.

Signature(s)  E6.

R.G.C. JENKINS & CO Date 12 February, 2004

12. Name, daytime telephone number and e-mail address of person to contact in the United Kingdom

D. C. Musker 020-7931-7141

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered YES in part 8, a Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- Part 7 should only be completed when a divisional application is being made under section 15(4), or when an application is being made under section 8(3), 12(6) or 37(4) following an entitlement dispute. By completing part 7 you are requesting that this application takes the same filing date as an earlier UK application. If you want the new application to have the same priority date(s) as the earlier application, you should also complete part 6 with the priority details.

Traffic Flow Determination in
Communications Networks

This invention relates to the field of communications networks. More
5 particularly, but not exclusively, it relates to a method of calculating traffic
paths in a communications network.

Background of the invention

There are different routing protocols used by network operators to
10 route traffic data in a communications network. Many network operators use,
for example, the interior gateway protocol (IGP) network routing protocols.

IGPs are a family of related protocols which have been in existence
for some time, and are widely deployed.

The family of IGPs includes different protocols, such as the Open
15 Shortest Path First (OSPF), the Intermediate System to Intermediate System
(IS-IS), the Routing Information Protocol (RIP), the Interior Gateway Routing
Protocol (IGRP) or the Enhanced Interior Gateway Routing Protocol
(EIGRP).

In IGPs the default routing strategy is shortest path first (SPF) routing.
20 Some IGPs allow link metrics to be assigned by the operator to each link in
the network. In these cases, a path is the shortest path in the sense that the
sum of link metrics along the path is minimal amongst all possible paths. In
other IGPs, the shortest path is simply the path with the lowest hop count (i.e.
the path with the shortest number of links). In all cases, IGPs use distributed

protocols so that, once the protocols are in a converged state, all nodes in the network have the same view of the shortest path between a given source and destination node pair.

5 Recently, the multi-protocol label-switching traffic engineering MPLS-TE protocol has been developed as an alternative routing protocol. However, the MPLS-TE is (as yet) much less widely used than the IGPs.

10 The method of routing traffic using the MPLS-TE protocol is considerably different from IGPs. The network operator may specify explicit path tunnels on the network. These paths carry the data traffic all the way from its source node to its destination node. The tunnel is configured by first informing the source node of the required tunnel configuration. Then, via MPLS-TE, sufficient details of the tunnel are communicated to all the nodes along the explicit path, so that these nodes can recognise packets in the tunnel and so that they know where to forward such packets.

15 When a data packet enters the network at the source node, the node selects for the packet one of the tunnels which starts at the source node and finishes at the destination node. The packet is then routed along the first link in this tunnel's path. When the packet arrives at a subsequent, intermediate path node before the destination, this intermediate node determines which
20 tunnel the packet is in. At such a node, the packet is then forwarded along the next explicitly defined link in its path.

In a **first scenario** where a network operator wishes to migrate from one routing protocol to another, the network operator may want to migrate

incrementally. For example, in migrating from an IGP routing protocol such as OSPF to the MPLS-TE protocol, a step-by-step approach is usually desired by network operators. This is due to the concern that an immediate switch to the new protocol may cause unpredicted disruption to the traffic on their network, and other unpredicted operational problems.

In addition to the migration scenario, a **second scenario** that causes difficulties for the network operators is when they wish to apply a routing protocol for the purposes of data collection, as some routing protocols have superior data collection characteristics to others.

As an example, network operators may wish to deploy MPLS TE for the purposes of gathering traffic matrix information about the size of the traffic between source-destination node pairs, while adhering to the IGP routing currently implemented on their network. This traffic matrix information may be used to improve the IGP routing metrics, or may be used as input to other network configuration or capacity planning decisions, or to assist in implementing an incremental migration strategy.

To clarify this point further, IGP metrics, or alternatively MPLS-TE paths, need to be calculated so as to avoid congestion. However this calculation depends on accurate end-to-end traffic matrix estimates. These estimates depend in turn on measurements that indicate how much traffic is passing between each pair of source and destination nodes. In this way the network operator can establish the effects of changing IGP link metrics, or of MPLS-TE paths for certain source-destination node pairs on various routes,

and can thus evaluate whether the network is susceptible to congestion. Installing IGP metrics or MPLS-TE paths without knowledge of the traffic matrix carries a high risk that congestion occurs in the network. If the network operator does not know the amount of traffic between the source and the destination nodes, he or she cannot predict congestion.

Once MPLS-TE paths are in place, it is usually possible to query each path to give an accurate measurement of the traffic flow over that path. However, nodes running IGP protocols do not provide this end-to-end traffic matrix information. In order to collect an accurate end-to-end traffic matrix from an IGP network in the absence of MPLS TE, resource-consuming methods need to be employed, such as the installation of costly hardware probes in the network or the installation of possibly disruptive analysis software.

Returning to the first migration scenario, in migrating IGP to MPLS-TE protocols, the network operator experiences a major difficulty unless a data collection process is implemented. This is because MPLS-TE paths that differ from the current routing cannot be calculated and safely installed without a traffic matrix, and the traffic matrix cannot be cheaply and accurately measured unless MPLS-TE paths are installed.

A **third scenario** is when network operators wish to deploy an alternative protocol due to its superior characteristics under network failure situations. As an example, network operators implementing IGP routing may wish to deploy MPLS TE because of its bandwidth protection capabilities which have

a high restoration speed, and can guarantee to protect quality of service temporarily after network link or node failures. However, they may wish to achieve this without affecting the network operation under normal conditions.

One aim of the present invention is therefore to alleviate the difficulties
5 of a migration between two different routing protocols.

Another aim of the present invention is to determine traffic paths in a communications network that mimic the routing behaviour under an existing or previously implemented routing protocol, whilst taking into account constraints imposed by another routing protocol.

10 Another aim of the present invention is to facilitate the extraction of traffic flow data from a communications network

Another aim of the present invention is to facilitate the provision of back-up facilities for operating a communications network

15 According to one aspect of the present invention, there is provided a method of determining traffic paths between one or more source-destination node pairs in a communications network, comprising starting from a first set of paths between said source-destination node pairs, determining a second set of paths between said source-destination node pairs while taking into account a set of constraints, such that said second set of paths emulates said first set of
20 paths.

According to another aspect of the present invention, there is provided a method of calculating traffic paths between one or more source-destination node pairs in a communications network, comprising starting from a first set

of paths between said source-destination node pairs, determining a second set of paths between said source-destination node pairs while taking into account a set of constraints, such that said second set of paths is similar to the first set of paths; and the traffic load of said second set of paths is similar to the traffic load of said first set of paths.

According to another aspect of the present invention, there is provided a method of calculating traffic paths between one or more source-destination node pairs in a communications network, comprising starting from a first set of paths between said source-destination node pairs determined using a first routing protocol, determining a second set of paths between said source-destination node pairs for use with a second routing protocol, such that the load balancing in said first and second routing protocols is similar.

Further aspects and advantages of the invention will be appreciated, by example only, from the following description and accompanying drawings, wherein

Figure 1 illustrates a simplified example of a communications network in which the present invention can be implemented;

Figure 2 illustrates the relationship of a network and a network management system, into which the present invention may be implemented;

Figure 3 is a diagram illustrating the method of determining traffic paths according to one embodiment of the present invention;

Figures 4A and 4B are diagrams illustrating individual steps of performing the method of Figure 3; and

Figure 5 is a diagram illustrating the method of determining traffic flow measurements according to a further embodiment of the present invention.

Figure 1 illustrates a simplified example of a communications network. Generally, such a network consists of nodes 2 and links 4. Links 4 carry data traffic between nodes 2. Here, link refers to the capability, as provided by various combination of equipment, to send traffic in one direction between two nodes. Traffic, usually in the form of data packets, enters the network at a source 6 node, is carried over the network via various links and nodes, and finally leaves the network at a destination node 8.

Network links have a numeric capacity indicating the magnitude of traffic they can carry. By altering their configuration, the nodes and links can be made to route traffic in different ways, i.e. to route traffic along different alternative path combinations. Network Quality of Service (QoS) is an important consideration when choosing the routing. If at any stage a link has more traffic on it than its capacity, it is said to be congested. Congested links lead to degraded quality of service. Therefore, network operators strive to configure the routing of their traffic carefully and try to avoid congestion. Other considerations regarding QoS are taken into account. They include path delay, that is, the length of time it takes for the traffic to travel from the source node to the destination node, and other QoS requirements such as jitter. In the following, congestion has been used to illustrate QoS

requirements, without loss of generality, and as an example only of the various QoS requirements.

Network routing is done via a routing protocol as described above.

The communications network 1 may be connected to other networks 10.

5 Figure 2 illustrates the relationship of the network and a network management system 60 in which the present invention may be implemented. Network management system 60 performs the network management functions for network 50. The network management system communicates with the network using a network management protocol, such as the Simple Network
10 Management Protocol (SNMP). The network management system 60 includes a processor 62 and a memory 64 and may comprise a commercially available server computer. A computer program performing the calculation of data traffic is stored in memory 64 and can be executed by processor 62.

15 **First Embodiment**

In the following, one embodiment of the present invention will be described. In this embodiment a network is operated using the IGP, and traffic paths are determined which facilitate migration to the MPLS-TE protocol, such that the implications on the over-all performance of the network are low.

20 Generally, IGP protocols may handle situations differently if there are multiple equally short paths between a given source and destination node pair. Some IGP protocols implement a so-called equal cost multiple paths (ECMP) extension. This is an example of a load balancing strategy. With these

protocols, when there are multiple shortest paths, all with the same metric sum, the nodes split the traffic (either actually or probabilistically) over the various outgoing links which are on these multiple paths.

Also, there exist various additional mechanisms to extend the basic
5 shortest path routing approach of IGPs, such as multiple routing areas (e.g. OSPF Areas, IS-IS levels), inbound route filtering or static routes. However, in all cases, given sufficient configuration information, for a given source destination pair, it is possible to determine which paths through the network the traffic will use, and which fraction of the end-to-end traffic will pass over
10 each link.

Also for the MPLS-TE protocol, different implementations are possible. Load sharing, for example, may be handled in different ways. If several tunnels are available that transmit data to a certain destination, the situation may be handled differently according to the network hardware and software involved
15 in implementing the tunnels. For example, the implementation of load sharing may depend on the hardware at the source node and/or the operating software which is present at the source node.

One example to illustrate is given in the following. Nodes manufactured by vendor A running a certain version of operating software
20 may, for example, allow equal splitting of traffic over several tunnels, and allow for a maximum of 16 tunnels only between a given source and destination pair. However, the network may use less than the maximum number of tunnels.

Alternatively, nodes manufactured by vendor B running different operating software may, for example, allow that each tunnel is assigned a so-called "load share" parameter which determines the fraction of the traffic which will use the tunnel. Also, the protocol may for example be
5 implemented such that splitting of traffic according to the load share parameter may be unequal as long as all traffic is allocated to some tunnel, and each tunnel is allocated a fraction of traffic which is a multiple of 1/16, and a maximum of 8 tunnels only may be used between a given source destination pair. Again, less than the maximum number of tunnels may be
10 used.

In summary, in MPLS-TE hardware and software limitations exist, which restrict which MPLS-TE tunnel combinations are possible between a source and destination node pair.

It is now referred to Figure 3. In order to determine traffic paths which
15 can be used in the MPLS-TE routing protocol, network topology data (step 101) and IGP route configuration data (step 102) are collected. In step 103 a set of source-destination node pairs are selected, for which traffic paths will be determined. This set of source-destination node pairs might be a subset of all such pairs in the communications network.

20 The data obtained in steps 101, 102 and 103 can now be used to build a data model of the IGP routing for each of the selected source-destination node pairs (step 104). For each pair, the IGP routing model determines the set of paths used to transmit data from the source node to the destination node in

the IGP routing and the fraction of traffic between the pair on each path. In step 105, MPLS-TE hardware and software limitations are specified, as these limitations constrain the determination of the traffic paths to be determined.

Using the IGP model for each source destination pair and taking into
 5 account the MPLS-TE hardware and software limitations, in steps 106 and 107 a set of MPLS-TE paths are determined for a particular source-destination pair. The paths are determined such that they are as closely as possible to the IGP paths for the particular node pair, while taking into account the constraints resulting from the MPLS-TE protocols and MPLS
 10 implementation.

Referring now to Figure 4, a simplified example will be presented to illustrate the embodiment described above.

Figure 4 illustrates a simplified network consisting of nodes A to H, wherein each link has the same link metric. Consider now a traffic flow
 15 between source node A and destination node D. It can be seen from Figure 4 that there are several shortest paths between nodes A and D, i.e. the five different paths $A \rightarrow B \rightarrow C \rightarrow D$, $A \rightarrow B \rightarrow F \rightarrow D$, $A \rightarrow E \rightarrow F \rightarrow D$, $A \rightarrow E \rightarrow H \rightarrow D$ and $A \rightarrow G \rightarrow H \rightarrow D$. If now ECMP load balancing is used, node A splits the traffic destined for node D equally over links $A \rightarrow B$, $A \rightarrow E$ and $A \rightarrow G$. So 1/3
 20 of the traffic from node A to node D passes over each of these links. Similarly, the flow is further split at later nodes, and at other nodes, such as node F, the traffic flow merges again before finally completely merging back again at the destination node D.

Listed below is the amount of traffic between all links of the communications network as determined from the IGP using ECMP load sharing:

	A→B:	1/3
5	A→E:	1/3
	A→G:	1/3
	B→C:	1/6
	B→F:	1/6
	E→F:	1/6
10	E→H:	1/6
	G→H:	1/3
	C→D:	1/6
	F→D:	1/3
	H→D:	1/2

15 Consider now that the MPLS-TE limitations demand that each path from node A to node D carries a fraction of traffic which is a multiple of 1/16.

If now the source-destination node pair A-D is selected, the IGP model depicted in Figure 4 A is used and the above described MPLS-TE limitation is taken into account, the traffic paths between nodes A and D are determined which “mimic” as closely as possible the IGP paths and the relative fraction of these paths as shown in Figure 4A.

The solution to this problem includes the same set of five paths
 $A \rightarrow B \rightarrow C \rightarrow D$, $A \rightarrow B \rightarrow F \rightarrow D$, $A \rightarrow E \rightarrow F \rightarrow D$, $A \rightarrow E \rightarrow H \rightarrow D$, $A \rightarrow G \rightarrow H \rightarrow D$,
 with the following traffic fractions:

	$A \rightarrow B \rightarrow C \rightarrow D$:	$3/16$
5	$A \rightarrow B \rightarrow F \rightarrow D$:	$2/16$
	$A \rightarrow E \rightarrow F \rightarrow D$:	$3/16$
	$A \rightarrow E \rightarrow H \rightarrow D$:	$3/16$
	$A \rightarrow G \rightarrow H \rightarrow D$:	$5/16$

This solution is illustrated in Figure 4 B.

10 Note that the total fraction of traffic placed on each link is as similar
 as possible to that of the IGP model as described above. For example on link
 $H \rightarrow D$, $\frac{1}{2}$ the traffic flows over this link in both cases. On link $F \rightarrow D$, on the
 other hand, it is not possible to match the IGP fraction exactly. The traffic
 paths resulting in a traffic flow which is closest to the IGP model results in a
 15 traffic flow of $5/16$ of the total traffic on this link. This is close to the $1/3$ of
 the traffic which is transmitted on this link in the IGP model.

Generally, for the embodiment described above with reference to
 Figure 3, the following quantities are used as an input to calculate paths
 taking into account MPLS-TE constraints which are “mimicking” or
 20 “emulating” the paths resulting from the use of IGP.

Input:

- For each node in the network, the MPLS-TE path constraints for that node
- A set of source-destination pairs
- 5 • For each source-destination pair
 - The set of all paths which is being used by the IGP protocol to route the traffic between the pair
 - For each link in the network, the fraction of the traffic between the pair which is being routed over that link

10 **Output:**

- For each source-destination pair
 - The MPLS-TE paths set, a subset of the IGP paths
 - For each MPLS-TE path, a fraction of traffic which should go over that path
- 15 such that
 - The MPLS -TE path constraints implied by the network hardware and software limitations are satisfied by the path set choice and path traffic fraction choices.
 - The MPLS-TE routing is as close to the IGP routing as possible in the sense that using the former routing leads to a fraction of end-to-end traffic on each link which is as close as possible to the fraction on the link when using the latter routing.
- 20

Considerations for problem decomposition

Generally, the above described problem can mostly be solved separately for each source destination pair. However, this separate treatment of the problem may introduce a systematic bias towards particular links or paths.

It is noted that in the example described above with reference to Figures 4 A and B, the link A→E receives 6/16 of the traffic in the “MPLS-TE case” (i.e. the determination as shown in Figure 4B), whereas the link receives 1/3 of the traffic in the IGP model. Thus, the same link carries slightly more traffic in the MPLS-TE case than in the IGP model. On the other hand, link A→B carries slightly less traffic in the MPLS-TE case compared to the IGP model (5/16 versus 1/3, respectively). Such discrepancies are unavoidable given the MPLS-TE limitations on tunnels between nodes A and D. However, if this bias in favour of placing traffic on link A→E instead of A→B were *systematic* across the MPLS-TE tunnels produced for many demands, the errors could multiply and so the routing could be much less similar to the IGP routing.

Such a systematic bias is prevented by randomly choosing between solutions which are equally close to the IGP routing, of which there are often several. In this way the discrepancies due to the MPLS-TE limitations at least partly cancel out when considering several source destination pairs.

Mathematical formulation

In the following, a mathematical formulation is given for the embodiment described above.

In view of the considerations regarding the problem consideration described above, in the mathematical model formulated below a single source-destination pair (s, t) is considered for which paths mimicking the IGP paths are computed.

The model below describes two particular MPLS-TE limitations: the constraints on the allowable load fractions and a maximum number of tunnels between a particular source-destination node pair. It is appreciated that these constraints serve merely as an example to illustrate the mathematical model, and that alternatively other limitations may be considered.

Constants

L is the set of all links. The letter i will range over links.

P is the set of all paths between s and t which are shortest paths in the IGP routing. The letter p will range over paths.

The 0/1 constants z_{ip} indicate whether link i is one of the links in path p .

The rational constants l_i ($0 \leq l_i \leq 1$) indicate the fraction of the traffic between the between s and t which passes over link i in the IGP routing.

F_s is a constant finite non-empty set of rational constants $\{f_1, f_2, \dots\}$: the allowed traffic fractions for MPLS-TE paths which start at node s . Zero must be a member of F_s . For each f in F_s , $0 \leq f \leq 1$.

M_s is a constant positive integer: the maximum number of MPLS-TE paths which can carry traffic from source node s to destination node t .

Variables

The 0/1 variables x_p indicate whether path p is one of the MPLS-TE paths.

The variables b_p range over the set F_s : the fraction of the traffic which is assigned to path p in the MPLS-TE routing.

Constraints

For any path p , it is used as an MPLS-TE path if and only if it has a non-zero fraction of traffic.

$$(\forall p \in P) x_p = 1 \leftrightarrow b_p \neq 0 \quad (C1)$$

The maximum paths limit at node s must be respected

$$\sum_{p \in P} x_p \leq M_s \quad (C2)$$

The sum of fractions of traffic over all MPLS-TE paths must be 1, i.e. the paths must carry all traffic.

$$\sum_{p \in P} b_p = 1 \quad (C3)$$

Optimisation function

The optimisation function is given by

$$\sum_{i \in L} |b_p z_{ip} - l_i| \quad (\text{Opt})$$

i.e. a minimisation of the sum over all links of absolute differences between the fraction of traffic between s and t on the link in the IGP routing and the fraction on the same link.

- 5 It is noted that alternatively other optimisation functions may be used. Generally, the optimisation functions relate the product $b_p z_{ip}$ (the MPLS-TE traffic fraction) to l_i (the IGP traffic fraction).

Search algorithm

- 10 In the following, a so-called “Generate and Test” search algorithm is presented, which outputs an optimal solution to this problem. It is noted that this algorithm avoids systematic link/path bias by breaking ties randomly, as described above.

1. *Solutions* := \emptyset
- 15 2. *BestCost* := ∞
3. Generate all member tuples of the cartesian product $F_s^{|P|}$
4. For each member b of $F_s^{|P|}$
 - a. For each tuple element b_p of b ,
 - i. If b_p is greater than 0 set x_p equal to 1
 - 20 ii. Else set x_p equal to 0
 - b. If constraints (C2) and (C3) hold for b and (Opt) for $b \leq$
BestCost

i. If $b < \text{BestCost}$ then $\text{Solutions} := \{b\}$ else $\text{Solutions} := \text{Solutions} \cup \{b\}$

ii. $\text{BestCost} := b$

5. If Solutions is non-empty, return a random member of Solutions .

5 It is noted that the above described algorithm relates to a single source destination pair. Usually, more than one source destination pair is considered in a network. In this case a single computer may execute the algorithm for all pairs. Alternatively, the computation can be distributed over a number of computers or nodes. Alternatively to the implementation as described above
10 with reference to Figure 2, the computation may be performed in one or more of the nodes 2 of the communications network as shown in Figure 1, or any other computer connected to the network.

The above embodiment may for example be implemented such that each node in the network is responsible for computing a subset of all pairs in a way
15 that all pairs are covered. For example a node might compute all pairs where it is the source. The fifth embodiment, which will be described below, relates to specific aspects of a distributed implementation.

Second Embodiment

20 The method described above may for example be used to temporarily move a part of the network traffic from an IGP routing protocol to a MPLS-TE routing protocol to collect accurate traffic flow data. As described above,

the MPLS-TE protocol provides for accurate measurements of the traffic flows.

This procedure will be explained with reference to Figure 5. Steps 101 to 107 of Figure 3 are equivalent to those in Figure 5, and are therefore
5 denoted with the same reference numerals.

In step 103, the network operator selects certain source-destination node pairs for which he wants to have traffic flow measurements. In steps 106 and 107 the MPLS-TE paths which are closest to the IGP paths used in the original network are calculated from the IGP model, taking into account the
10 MPLS-TE constraints as described earlier with reference to Figure 3.

In step 108, the calculated MPLS-TE paths are installed on the network and the network is operated by using the MPLS-TE routing protocol for part of the network, i.e. the source-destination node pairs selected in step 103. Thus, the network operator receives accurate end-to-end traffic
15 measurements for the selected part of the network in step 109.

Although the paths for the network operating with IGP routing protocol may be slightly different to the paths used for MPLS-TE routing, the MPLS-TE routing will have very similar properties to the IGP routing, and in particular the resulting link loads will only be slightly affected by the
20 introduction of the MPLS-TE routing.

If the MPLS-TE tunnels are no longer required after the end-to-end traffic data has been collected, they can be removed, and the network is operated by reverting to IGP routing.

Third Embodiment

The method described above may also be used if a network operator wants to decommission permanently the IGP protocol in favour of MPLS-TE technology, for example because the network operator considers the MPLS-TE protocol to offer better manageability and flexibility compared to the IGP.

In this case the network operator applies the migration process as described above with reference to Figures 3 and 5 to all source-destination pairs, and in this way collects traffic data for each MPLS-TE tunnel. With these traffic data the network operator can set the resource reservation setup protocol (RSVP) bandwidths of the tunnels according to the traffic flow measurement obtained to ensure smooth running of the communications network after migration to the new routing protocol.

In this way the “risk” of migration from a network operating with an IGP to a network operating using the MPLS-TE protocol is reduced as the network operator knows that the traffic in the MPLS network is routed in paths very similar to those previously used. Therefore, if congestion was absent in the IGP routing, it should also be absent in the IGP-mimicking MPLS-TE routing.

Fourth Embodiment

A network operator may want to migrate some of the network traffic from IGP routing to MPLS-TE routing in order to provide back-up facilities such as secondary paths in case of a network failure. Such back-up facilities
5 are for example provided by a MPLS network in form of the so-called FastReRoute (FRR) facility bypass tunnels. These mechanisms provide back-up facilities with fast failover times and can form part of a bandwidth protection mechanism.

In this case the network operator applies the general migration process
10 to all source-destination pairs, and collects traffic data to determine the link load on each link. In the next step the network operator assigns a primary bandwidth to each link.

The primary bandwidth of each link is assigned such that there will be no congestion under failure, as long as the link load is below the assigned
15 amount when the failure occurs. The bandwidth to be assigned is calculated on the basis of the measured link load. The network operator may, for example, select the maximum link load observed over a given time period as the assigned primary bandwidth for each link. The network operator also assigns a secondary or backup bandwidth, which specifies the amount of
20 traffic which can be carried by the link in addition to the primary bandwidth in case of a network failure.

The secondary bandwidth is assigned such that no unacceptable congestion on the link is caused. The network operator may, for example,

select a certain percentage of the link capacity minus the primary bandwidth as the secondary bandwidth. Subsequently, the secondary tunnels are calculated to ensure there is no congestion, without having knowledge of the size of end-to-end traffic flows.

5 It is noted that the secondary tunnels may not use the FRR load balancing features.

Then the secondary tunnels can be installed in the network.

Fifth embodiment

10 It may be the nodes in the network that create and maintain the MPLS-TE paths that mimic IGP routing. Individual nodes could be instructed to maintain a set of source destination pairs, for example all those pairs from the individual node to a set of selected destination nodes. They would then create the paths as described above, and update them whenever there were changes
15 in the topology. The resulting communications network is routed using MPLS-TE but the routes match IGP routing closely, without the management overhead associated with MPLS-TE.

In order to avoid a systematic bias in the determination of the traffic paths, a documented pseudo-random method may be used to select an
20 “optimal” solution, so that an external tool could simulate the node choices.

There would be time-lag parameter, so that routes would not be adjusted immediately after any topology change. Setting different time-lags on different nodes would ensure a smooth transition from old to new routes whenever the topology changed.

- 5 The benefits of a distributed implementation for the network operator are enhanced simplicity and reliability.

Sixth embodiment

10 It is very often the case that there are a small number of source-destination pairs contribute the majority of all traffic. The network operator may wish to take advantage of this by only converting those source-destination pairs to MPLS-TE, using routes that emulate IGP routing. Alternatively or in addition, the network operators may convert only those pairs that, according to IGP routing, place traffic on the most congested links.

- 15 Having converted the selected (high traffic or congestion) source-destination node pairs, measured the traffic flow as described above and confirmed that the majority of traffic carried by the most congested links is carried within tunnels, the operator may then convert the source-destination pairs carrying smaller traffic flows for example by selecting an "arbitrary"
- 20 shortest path, without necessarily trying to emulate IGP load balancing, as the network operator knows that these flows are either small or do not use the most congested links.

One example method of selecting an initial set of “large flows” without the network operator knowing the traffic matrix, is for the operator to use traffic flow analysis and IGP route simulation. Based on link load measurements, the network operator selects traffic flows for which the traffic flow analysis determines that they are not larger than a certain predetermined threshold. Based on IGP route simulation, the network operator selects those pairs amongst them that are routed over the most congested links.

While it is often the case that network operators partition their nodes into those that are sources and destinations and those that are mid points where no traffic originates, the network operator may alternatively choose to have nodes performing both functions. In the case the traffic flow figures obtained for tunnels between such a dual-purpose source node and the destination node may include any traffic not already carried by tunnels, that was going to the same destination and traversed that node on the way.

In this case an operator wishing to only operate part of their network using MPLS-TE needs to take care to ensure that the only traffic being carried by the tunnels originates at those nodes. This can be done in various ways. For example, the operator may force traffic to avoid an existing tunnel by using static routes, or by creating additional tunnels which ensure that flows from different sources are carried in disjoint tunnels.

Whilst in the above described embodiments an optimal search algorithm has been described, it is appreciated that alternatively so-called heuristic algorithms may be used, which might not return an optimal solution.

5 Whilst in the above described embodiments the "Generate and Test" approach has been described, it is appreciated that other search algorithms or algorithm engineering techniques can be applied alternatively, as for example constraint programming, mathematical programming, operations research, artificial intelligence.

10 Whilst in the above described embodiments the migration from a IGP network to an MPLS-TE network has been described, it is appreciated that the method can alternatively used for migrating networks between other routing protocols.

15 It is to be understood that the embodiments described above are preferred embodiments only. Namely, various features may be omitted, modified or substituted by equivalents without departing from the scope of the present invention, which is defined in the accompanying claims.

CLAIMS:

1. A method of determining traffic paths between one or more source-destination node pairs in a communications network, comprising
5 starting from a first set of paths between said source-destination node pairs, determining a second set of paths between said source-destination node pairs while taking into account a set of constraints, such that said second set of paths emulates said first set of paths.
- 10 2. A method according to claim 1, wherein the second set of paths is determined such that the traffic load on said second set of paths emulates the traffic load on said first set of paths.
- 15 3. A method according to claim 1 or 2, wherein the first set of paths are included in a routing and load model for said source-destination node pairs related to a first routing protocol.
- 20 4. A method according to claim 3, wherein said routing and load model takes into account the network topology, the route configuration resulting from the use of the first routing protocol and/or a selection of source-destination node pairs.

5. A method according to any preceding claim, where the first set of paths is related to the use of a first routing protocol.

6. A method according to any preceding claim, where the second set of paths is determined for use with a second routing protocol.

7. A method according to any preceding claim, wherein the second set of paths is determined such that the routing using a second routing protocol is similar to the routing using a first routing protocol

10

8. A method according to any preceding claim, wherein said set of constraints is related to the second set of paths.

9. A method according to any preceding claim, wherein said constraints result from network nodes limitations and/or routing protocol constraints related to said second set of paths.

15

10. A method according to any of claims 4 to 9, wherein said first routing protocol includes an interior gateway protocol.

20

11. A method according to any of claims 4 to 10, wherein said first and/or said second routing protocol applies load balancing.

12. A method according to any of claims 4 to 11, wherein said first routing protocol includes an equal cost multiple paths extension.

13. A method according to any of claims 5 to 10, wherein in said
5 second routing protocol data are routed on pre-determined tunnels.

14. A method according to any of claims 5 to 11, wherein said
second routing protocol includes a multi-protocol label-switching traffic
engineering protocol.

10

15. A method according to any preceding claim, wherein said
constraints comprise a maximum number of paths between each source-
destination node pair.

15 16. A method according to any preceding claim, wherein said
constraints comprise that the traffic between a particular source-destinations
node pair is load-balanced such that the share of traffic along any paths is a
fraction with constrained integer numerator and denominator.

20 17. A method according to any preceding claim, wherein a search
technique is used to determine said second set of paths.

18. A method according to any preceding claim, wherein one of the following search techniques are used to determined said second set of paths: "generate and test" search algorithm, constraint programming and/or mathematical programming.

5

19. A method according to any preceding claim, wherein an optimal search algorithm is used.

20. A method according to any of claims 1 to 18, wherein a heuristic search algorithm is used.

10

21. A method according to any preceding claim, wherein each source-destination node pair is treated independently.

22. A method according to claim 21, comprising a method of avoiding a systematic bias for particular paths.

15

23. A method according to claim 22, wherein ties between symmetric solutions are broken randomly.

20

24. A method of calculating traffic paths between one or more source-destination node pairs in a communications network, comprising

starting from a first set of paths between said source-destination node pairs, determining a second set of paths between said source-destination node pairs while taking into account a set of constraints, such that

- i) said second set of paths is similar to the first set of paths; and
- 5 ii) the traffic load of said second set of paths is similar to the traffic load of said first set of paths.

25. A method of calculating traffic paths between one or more source-destination node pairs in a communications network, comprising

- 10 starting from a first set of paths between said source-destination node pairs determined using a first routing protocol, determining a second set of paths between said source-destination node pairs for use with a second routing protocol, such that the load balancing in said first and second routing protocols is similar.

15

26. A method of operating a communications network, comprising switching at least some network traffic from a first routing protocol to a second routing protocol, wherein the method includes a method of calculating traffic paths according to any of claims 1 to 25.

20

27. A method of measuring traffic between a plurality of source and destination nodes in a communications network comprising the method according to any of claims 1 to 26.

28. A method according to claim 26 or 27, wherein at least some of the traffic is protected using secondary tunnels.

5 29. A method of providing secondary paths for a communications network, comprising the method of any of claims 1 to 26.

30. A method according to claim 29, wherein a measured maximum link load is used as the primary bandwidth for each link.

10

31. A method according to claim 29 or 30, wherein part of the remaining link capacity is used for the secondary paths.

32. A method according to claim 29, 30 or 31, wherein the secondary paths are determined for the non-load balanced case.

15

33. An apparatus for calculating paths in a communications network, adapted to perform the method of any preceding claim.

20 34. An apparatus according to claim 33, wherein said apparatus includes one or more nodes of said communications network.

35. An apparatus according to claim 34, wherein said one or more nodes update the calculation of traffic paths after a change in the network topology.

5 36. An apparatus according to claim 35, wherein a time lag is introduced between the change in the network topology and an update of the traffic paths.

10 37. A network management system for managing a network, adapted to perform the method of any of claims 1 to 32.

38. A computer program for performing the method of any of claims 1 to 32 when operated in a computer system.

15 39. A method of calculating paths in a communication network as described herein with reference to the accompanying drawings.

ABSTRACT**Traffic Flow Determination in
Communications Networks**

5

A method of determining traffic paths between one or more source-destination node pairs in a communications network, comprising starting from a first set of paths between said source-destination node pairs, determining a second set of paths between said source-destination node pairs while taking
10 into account a set of constraints, such that said second set of paths emulates said first set of paths.




Fig. 1




Fig. 2

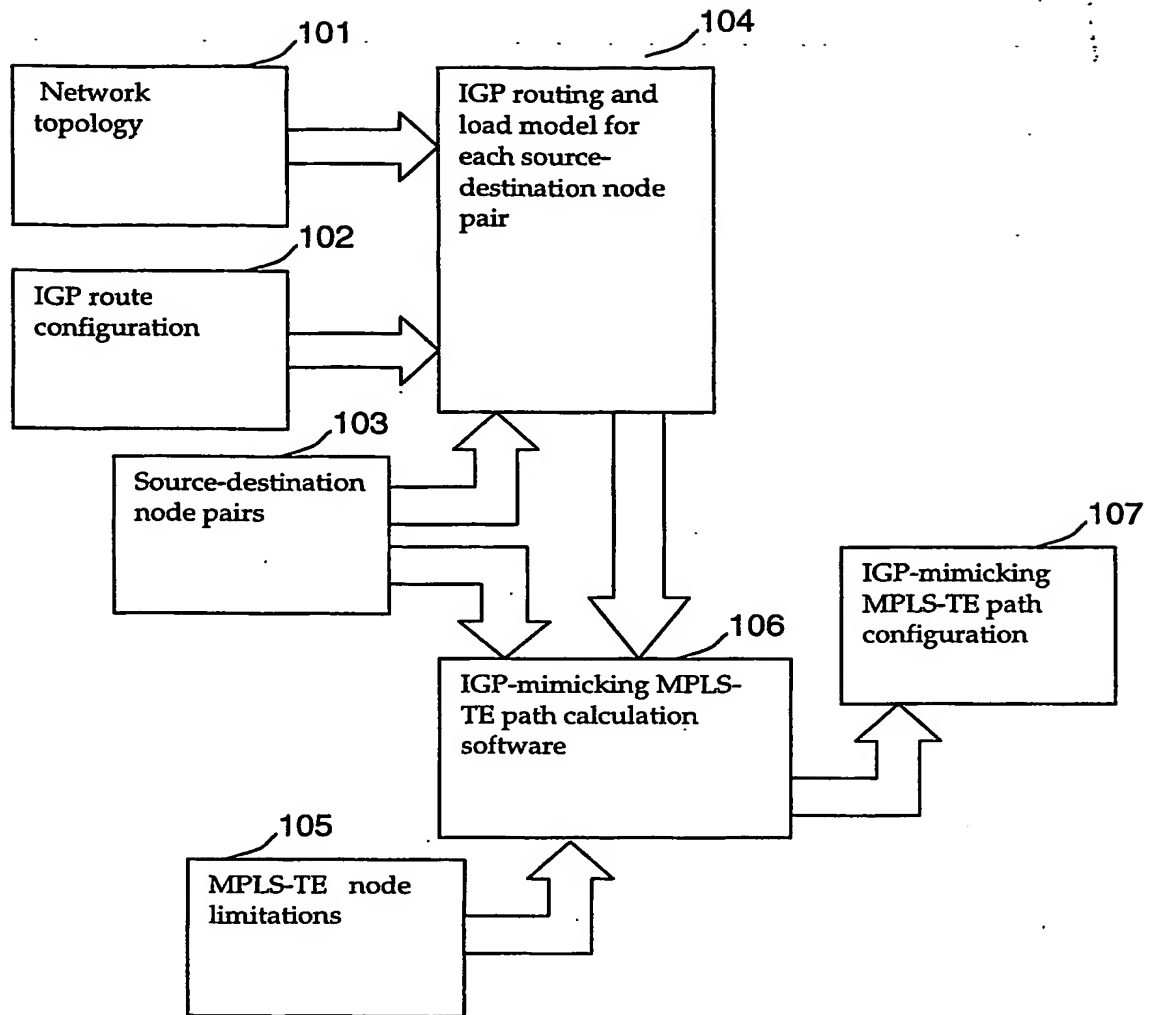


Fig. 3

4/5




Fig. 4A




Fig. 4B

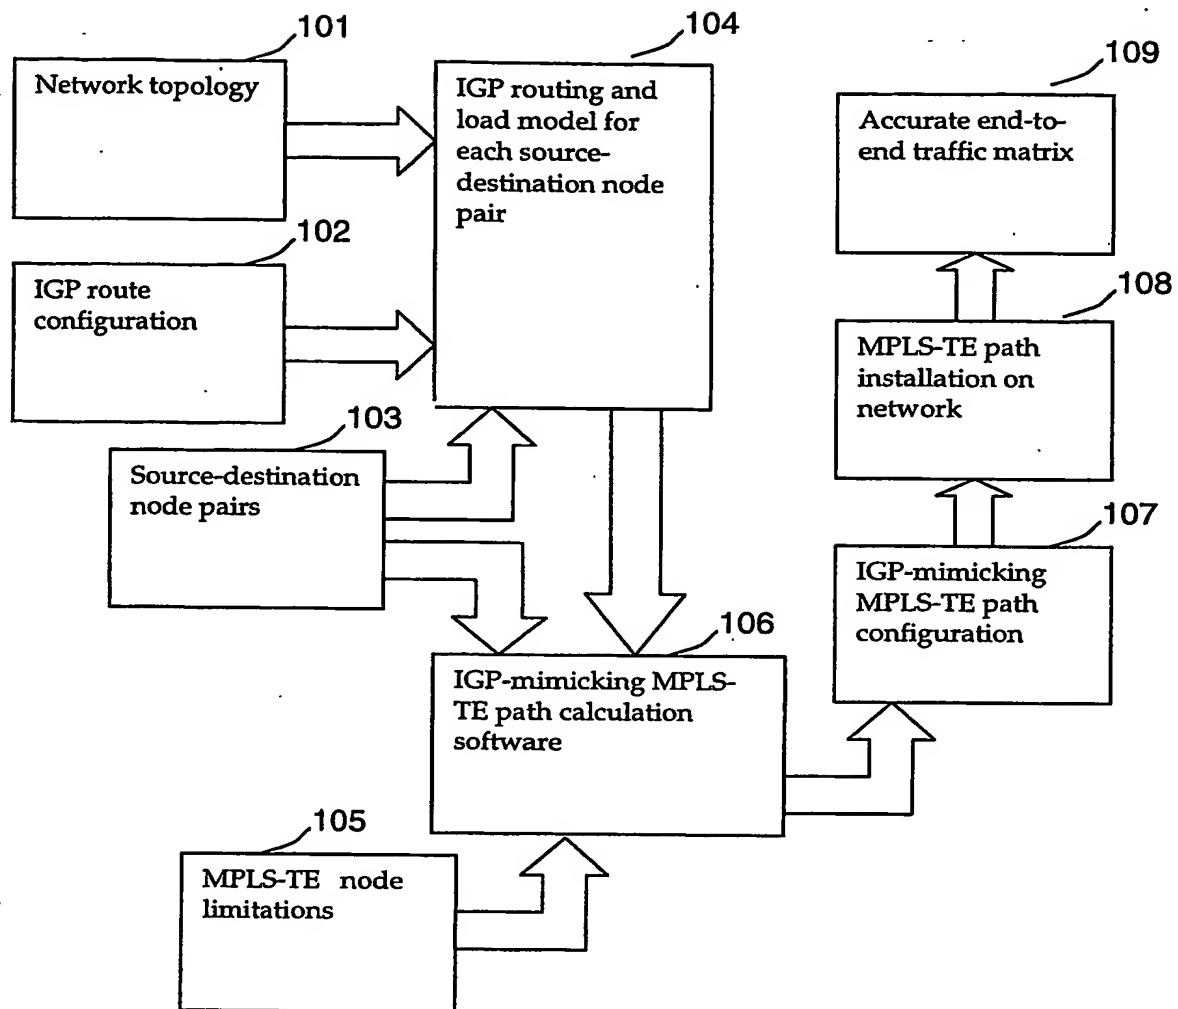


Fig. 5

Document made available under the Patent Cooperation Treaty (PCT)

International application number: PCT/EP05/001479

International filing date: 14 February 2005 (14.02.2005)

Document type: Certified copy of priority document

Document details: Country/Office: GB
Number: 0403159.7
Filing date: 12 February 2004 (12.02.2004)

Date of receipt at the International Bureau: 07 April 2005 (07.04.2005)

Remark: Priority document submitted or transmitted to the International Bureau in compliance with Rule 17.1(a) or (b)



World Intellectual Property Organization (WIPO) - Geneva, Switzerland
Organisation Mondiale de la Propriété Intellectuelle (OMPI) - Genève, Suisse